Numerical Weather Prediction

The MOGREPS short-range ensemble prediction system Part II: Performance and Verification

NWP Technical Report No. 498

Neill E. Bowler, Alberto Arribas, Kenneth R. Mylne,

Kelvyn B. Robertson, Sarah E. Beare & Timothy P. Legg

email:nwp_publicationsmetoffice.gov.uk

© Crown copyright

The MOGREPS short-range ensemble prediction system Part II: Performance and Verification

Neill E. Bowler* Alberto Arribas, Kenneth R. Mylne, Kelvyn B. Robertson, Sarah E. Beare & Timothy P. Legg

Met Office, FitzRoy Road, Exeter, EX1 3PB, UK.

Abstract:

The Met Office has recently introduced a short-range ensemble prediction system known as MOGREPS, which consists of a high-resolution regional ensemble nested within a global ensemble. A full description of the system is included in part I of this paper. MOGREPS has been trialled within the operational suite at the Met Office since August 2005. Here we describe the performance and routine verification of the system. Results show a good performance, with the regional ensemble better than the global ensemble and the ECMWF ensemble (as measured by the Brier Skill Score for forecasts of temperature greater than 10^{o}C and 15^{o}C, and wind speed greater than force 5).

Crown Copyright © 2007

KEY WORDS ensemble forecasting short-range NWP

*Correspondence to: N. E. Bowler, Met Office, FitzRoy Road, Exeter, EX1 3PB, UK. E-mail: Neill.Bowler@metoffice.gov.uk

Crown Copyright © 2007

1 Introduction

MOGREPS is the new short-range ensemble prediction system developed at the Met Office. It consists of a global and regional ensemble, with the global providing the lateral boundary conditions and initial condition perturbations for the regional. The ensemble transform Kalman filter (ETKF) is used to calculate the perturbations to the initial condition and two schemes (the Stochastic Convective Vorticity scheme, SCV, and the Random Parameters scheme, RP) are included to represent structural and sub-grid-scale sources of model uncertainty. The system and its setup are fully described in part I of this paper (Bowler *et al.*, 2007).

Here we focus on the performance and verification of the system. The complete MOGREPS (including both regional and global ensembles) has been running quasi-operationally since August 2005. Being a new system, the verification software is still under development and the verification at this stage is limited.

The paper is structured as follows: an example forecast - demonstrating significant short-range ensemble spread - is shown in section 2. Section 3 analyses the interaction between the perturbations to the initial condition and the stochastic physics. Analysis of the global ensemble performance is presented in section 4. Section 5 presents the verification of the regional ensemble. Conclusions are drawn in section 6.

2 Example forecast

On the evening of 2 December 2005 a violent storm was observed in the English channel, with maximum observed winds of Beaufort force 11 (56 knots). This case illustrates the benefits of having a high-resolution short-range forecasting system. T+24h forecasts for the MOGREPS NAE ensemble, valid at 18Z on 2 December 2005 are shown in figures 1 and 2. Figure 1 shows the forecast of PMSL from the MOGREPS regional ensemble. The large-scale differences between the ensemble members at this lead time is small, with all the ensemble members having a similar structure of the cyclone. However there are important differences in detail. The lowest closed contour varies between 972 hPa for six of the members and 964 hPa for five of the members. Figure 2 shows the wind speed on model level 6 (approximately 610m above surface) which is used by forecasters as an estimate of the likely strength of a wind gust. Member 22 indicates gusts in the western part of the channel in excess of 70 knots, which is consistent with what was observed. Members 8 & 14 keep the strongest winds further to the south, suggesting much lower wind speeds within the channel. Feedback from forecasters indicated that this alert from MOGREPS aided them in issuing a more severe (and accurate) gale warning to shipping.

3 Interaction between initial condition and model uncertainty perturbations

Since the MOGREPS system uses a filtering method to provide the initial condition perturbations, model perturbations are able to feed back to the initial condition over a series of cycles. To analyse the effect of this interaction, an experiment was performed comparing the growth of the stochastic physics perturbations with the behaviour of perturbations to the initial condition. Three short-range (3-day forecast)

Crown Copyright © 2007

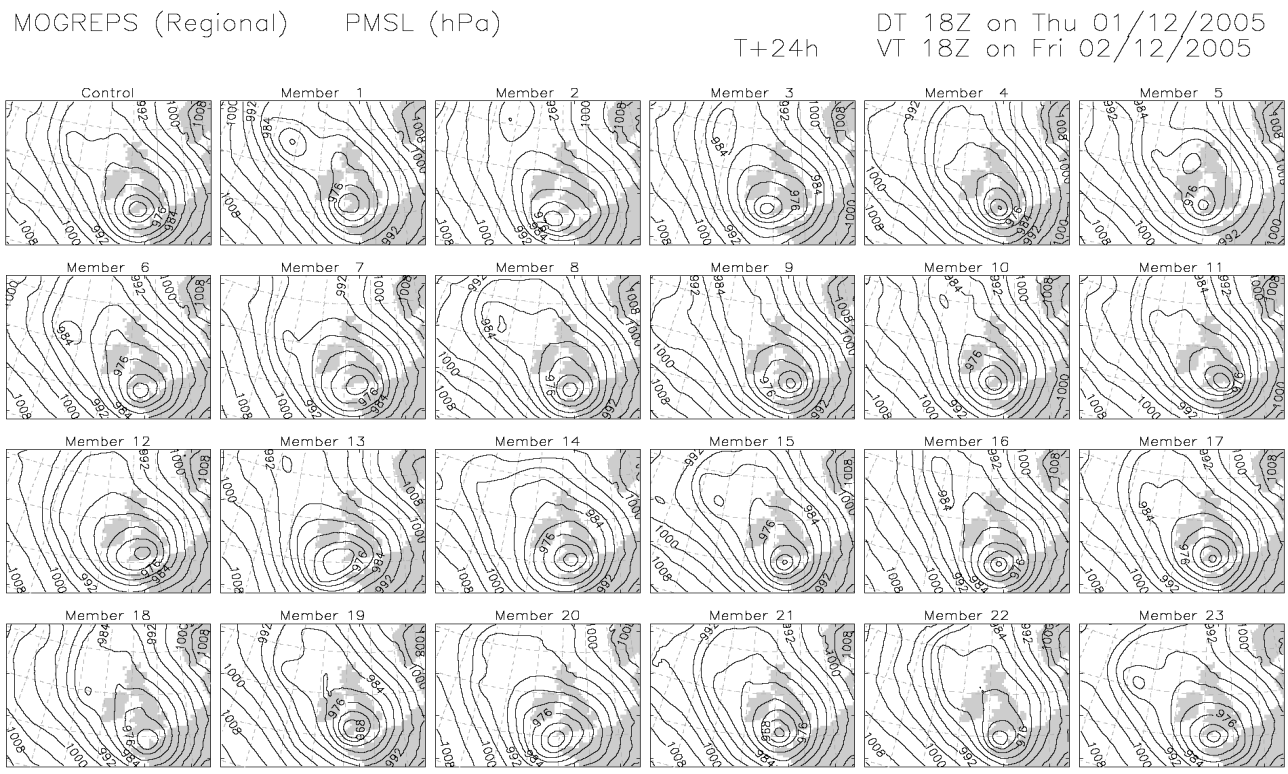

Figure 1. PMSL postage stamps from MOGREPS regional ensemble

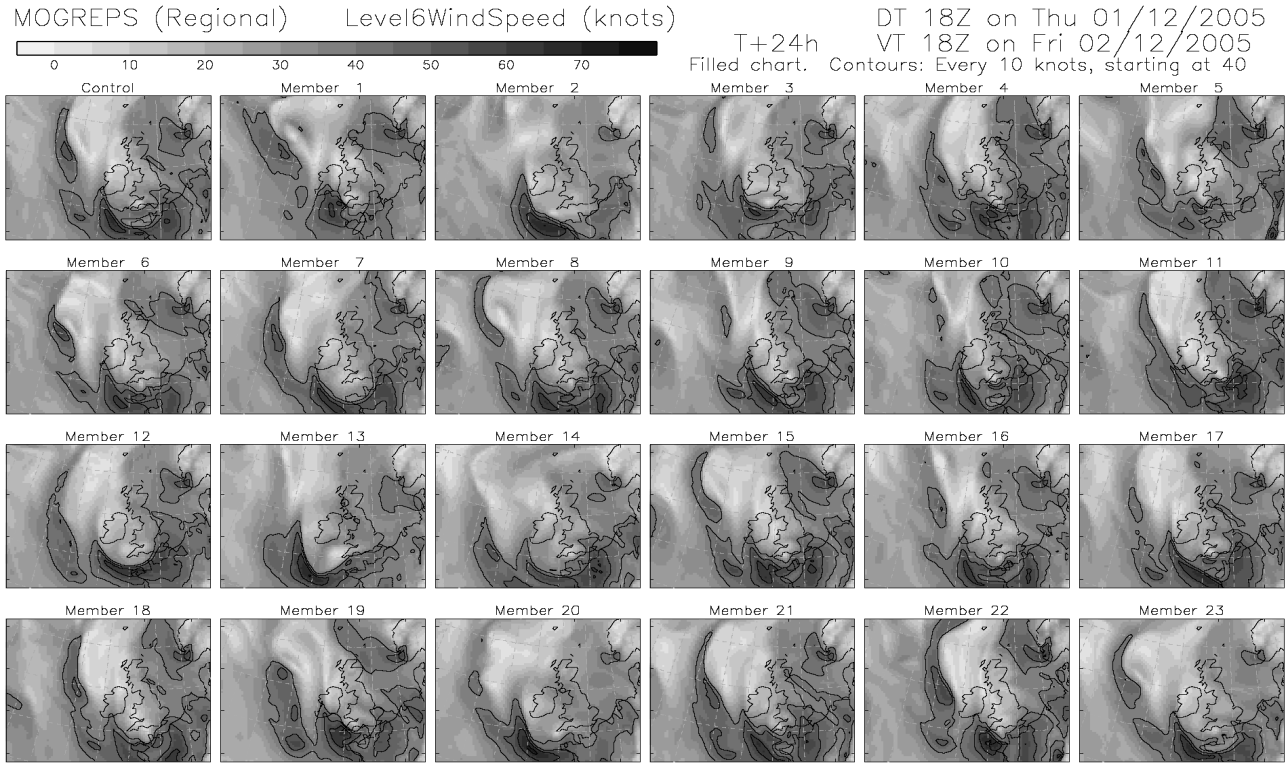

Figure 2. Level 6 wind speed postage stamps from MOGREPS regional ensemble

8-member ensembles with different configurations were run on 29 January 2004. The first one (IC, from initial condition), only

Crown Copyright © 2007

included perturbations to the initial condition, calculated using the Ensemble Transform Kalman Filter as described in part I (Bowler *et al.*, 2007). The second one (ME, from Model Error) included both stochastic physics schemes but no perturbations to the initial condition. The third one (IC+ME) included perturbations to the initial condition and the stochastic physics schemes. The initial condition perturbations were generated using a 7-day spin-up cycle of the ETKF, without stochastic physics. The same initial condition perturbations were used for the IC and IC+ME runs.

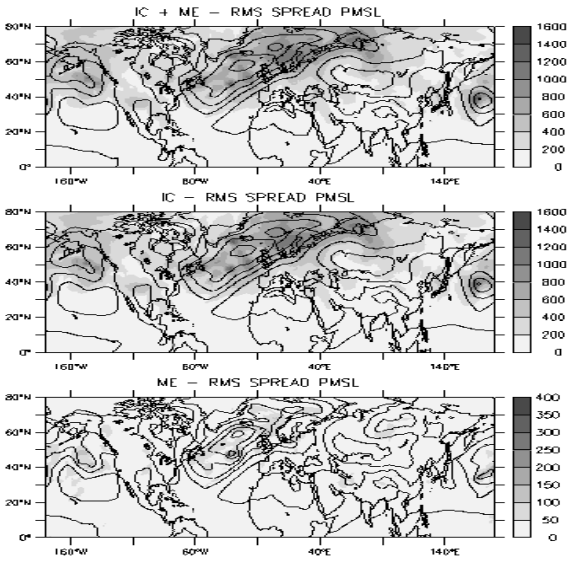

Figure 3. Spatial distribution of spread for PMSL after 72-h (shaded) and ensemble mean values (contour). Three ensemble configurations are shown: perturbations to initial condition and model error (IC+ME, top), only perturbing the initial condition (IC, centre) and only model error (ME, bottom).

The spatial distribution of PMSL spread after 72 hours is shown in Figure 3 for all configurations: IC+ME (above), IC (centre) and ME (bottom). As can be seen, in all three cases the largest values of spread are found over the same regions: North Atlantic-Scandinavia, east of the Japanese coast and near the North American west coast. This case illustrates that the spatial distribution of spread does not necessarily depend on the mechanism used to introduce the perturbations but on the background flow. This implies that the stochastic physics scheme and the initial condition perturbations could be exploring common areas of the phase space. If we compare the spread in the IC and ME ensembles, it seems that the ME spread is approximately 13% of the IC spread. When perturbations to the initial condition and the stochastic physics scheme are both included, as in IC+ME, the amount of spread is only increases by 5% with respect to the IC ensemble. This can be observed in Figure 4, where the time evolution of the average spread over the northern hemisphere is shown.

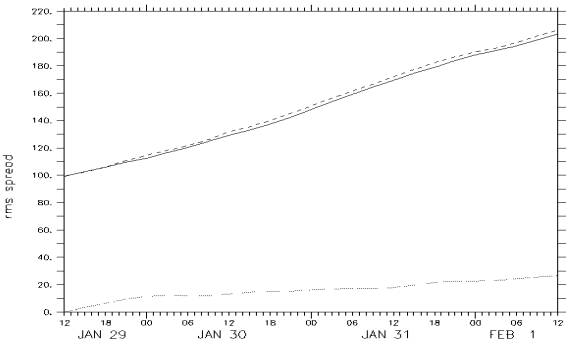

Figure 4. Time-evolution of PMSL spread for three different ensemble configurations: IC+ME (dashed line), IC (solid line) and ME (dotted line).

The role of the stochastic physics schemes representing sub-grid sources of model uncertainty fits well with the Tribbia and Baumhefner (2004), view of the energy cascade. They argued that the primary role of the inverse cascade of energy from small to large scales is to seed disturbances in the baroclinically active region. Once there, the perturbations would organise within the synoptic structures and amplify, extracting energy from the large-scale background flow, not the smaller scales. A consequence

Crown Copyright © 2007

of this is that it is not possible to design schemes to represent model uncertainties fully orthogonal to the initial condition uncertainties, which helps to explain the overlap in spread between IC and ME experiments.

4 Broad-scale verification of the global ensemble

In part I of this paper, the average spread of perturbations to the 500 hPa height field was shown. Here figure 5 shows the average ensemble spread and RMSE of the ensemble mean forecast for T+72h forecasts of 500hPa height valid during February 2006. An enlarged version of this data for the NAE domain is shown in figure 6. The RMSE is calculated relative to the analyses valid at the same time. A number of features are evident. Firstly, the broad-scale pattern of the spread matches that of the RMSE, though some of the details differ. Secondly, the spread seems to have a smoother spatial variation than the RMSE, this is likely to be due to extra sampling variation in the RMSE field, since the spread is based on 23 ensemble perturbations, and the RMSE being derived from a single verification field.

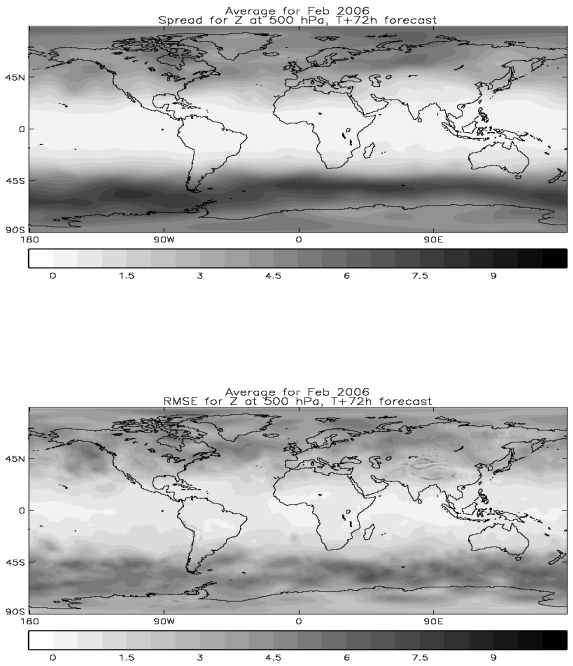

Figure 5. Average forecast spread (top) and RMSE (bottom) of the ensemble mean forecast at T+72 hours for the 500 hPa height field (in dm) for February 2006.

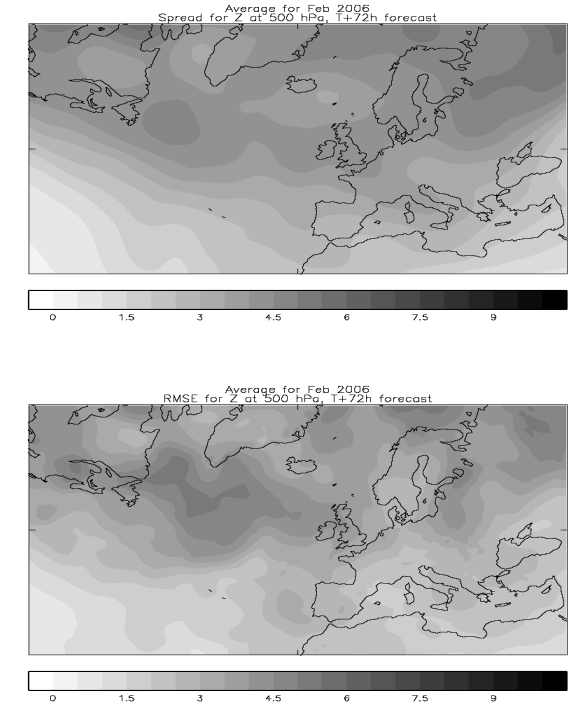

Figure 6. Average forecast spread (top) and RMSE (bottom) of the ensemble mean forecast at T+72 hours for the 500 hPa height field (in dm) for February 2006.

Figure 7 shows the same data as figure 5, but the spread and RMSE have been binned according to latitude. This indicates that the spread is too small in the tropics, and too large near the north pole and in the southern extra-tropics. The normalisation of the spread of the initial condition perturbations is performed by the ETKF inflation factor. This uses all radio-sonde observations outside the tropics and ensures that at T+12 the forecast spread matches the RMSE of the ensemble mean. If the ensemble spread grows more rapidly in one

Crown Copyright © 2007

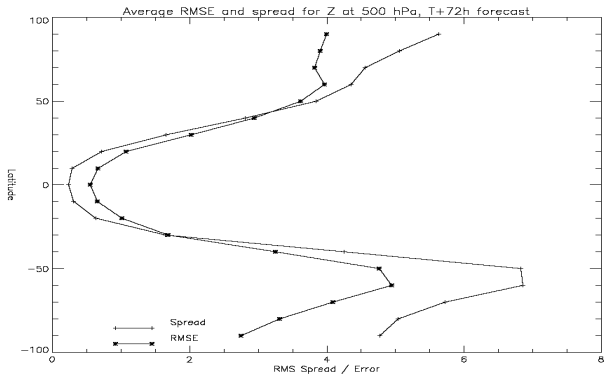

Figure 7. Variation of the average spread and RMSE for ensemble forecasts of 500 hPa height plotted against latitude at a lead time of T+72 hours for Jan 2006.

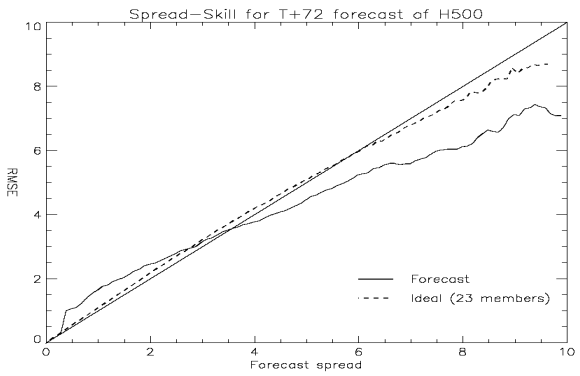

Figure 8. Relationship between the forecast spread and RMSE of the ensemble mean for forecasts of 500 hPa height at a lead time of T+72 hours. The result that would be seen for a perfect ensemble with 23 members, for this distribution of forecast spread, is shown by the dashed line.

region than another, then the ensemble initial condition perturbations will be larger in the region with the faster growth. The mixing of the ETKF perturbations through the transform matrix will attempt to offset this effect, and ensure that the amplitude of the initial condition perturbations is consistent with estimate of the error at all locations. However, this mixing limited by the fact that the MOGREPS ensemble only has 23 perturbed members, which is clearly insufficient to correct for the differences in the growth rate of the perturbations. A development of MOGREPS is planned to use the local ETKF which was proposed by Bowler (2006a) to determine the initial condition perturbations, which should alleviate this problem.

The average RMSE of the 72h ensemble mean forecast of 500 hPa height for the global model, binned according to forecast spread, is shown in figure 8. The verification is performed against analysis over the NAE domain, and is for forecasts valid between 17th October 2005 and 9th January 2006. There is clearly a good relationship between the spread and the RMSE. The dashed line shows what would be achieved by a perfect ensemble with 23 members. The perfect ensemble curve lies away from the diagonal because of the limited ensemble size, as was noted by Roulston (2005). Some of the spread-error relationship must be attributed to the latitudinal variation of the spread and error over the NAE domain (see Fig 6).

Figure 9 shows the rank histogram (Anderson, 1996; Hamill and Colucci, 1997) of T+72 forecast of 500hPa height verified in the same way as Fig 8. The curve is close to the ideal flat line, with only a slightly higher number of outliers than the ideal case. The frequency of outliers is 10.0%, which is 1.20 times larger than the frequency of outliers that would be seen for a perfect ensemble. In figure 9 of Buizza *et al.* (2005) the frequency of outliers relative to the expected frequency is reported for the Canadian, NCEP and ECMWF ensemble systems. If we assume that our ensemble would give 1.2 times the expected number of outliers in the configuration used for the test in that paper, then it can be calculated that the level of excess outliers in the MOGREPS ensemble would correspond to a value of 3.6%. This is closer to zero than the value achieved by the NCEP and ECMWF ensembles at 3 days lead time.

Crown Copyright © 2007

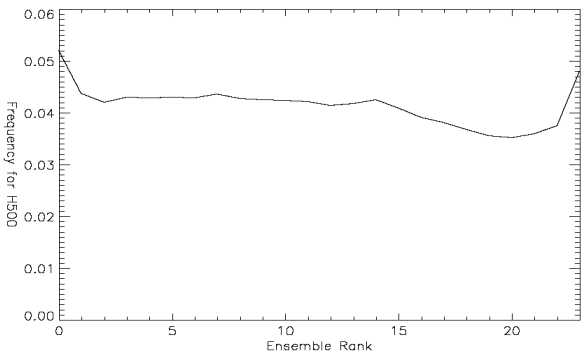

Figure 9. Rank histogram for forecasts of 500 hPa height, at a lead time of T+72 hours.

5 Verification of the regional ensemble

The following verification is for probability forecasts, mainly from the NAE ensemble. The variables verified are forecasts of screen temperature and precipitation, and are verified against observations. These are aimed at providing a user-oriented verification since surface variables constitute the main products of the MOGREPS system. We are aware that account should ideally be taken of the effect of observation errors (Bowler, 2006b) and varying climatology (Hamill & Juras, 2007), but this is currently beyond the capability of our verification system.

The Brier skill score (BSS), reliability score and resolution score are used to assess forecast quality. These scores derive from the well-known Brier score (Wilks, 2006, p286)

$$BS = \frac{1}{N}\sum_{c=1}^{N_c} F_c\left(\frac{O_c}{F_c} - p_c\right)^2 - \frac{1}{N}\sum_{c=1}^{N_c} F_c\left(\frac{O_c}{F_c} - f\right)^2 + f(1-f) \quad (1)$$

f is the sample climatological value of the probability of the event, that is

$$f = \frac{\sum_{c=1}^{N_c} O_c}{\sum_{c=1}^{N_c} F_c}. \quad (2)$$

O_c is the number of times the event was observed to occur in the probability category c, F_c is the number of forecasts of the event in the same category, N_c is the number of categories, N is the total number of forecasts and p_c is the forecast probability. The BSS is derived from the Brier score as

$$BSS = 1 - \frac{BS}{BS_{ref}}. \quad (3)$$

The in-sample climatology is usually taken as the reference forecast, which gives

$$B_{skill} = B_{Resolution} + B_{Reliability} - 1 \quad (4)$$

where

$$B_{Resolution} = \frac{1}{Nf(1-f)}\sum_{c=1}^{N_c} F_c\left(\frac{O_c}{F_c} - f\right)^2 \quad (5)$$

measures the propensity of the forecast to give high or low values of the probability (as opposed to forecasting climatological values) and

$$B_{Reliability} = 1 - \frac{1}{Nf(1-f)}\sum_{c=1}^{N_c} F_c\left(\frac{O_c}{F_c} - p_c\right)^2 \quad (6)$$

is a skill score measuring the reliability of the forecast. Ideally, $p_c = \frac{O_c}{F_c}$ and the reliability score would be one.

Verification of the forecasts from the MOGREPS global and regional ensembles, and the ECMWF ensemble for the months of March-May 2006 are shown in figures 11 to 16. The ECMWF ensemble is used as a comparison since it is the EPS currently available to the Met Office. If MOGREPS performance is inferior to it, then there is little justification for the extra computational expense. The verification is performed against surface observations for 56 stations in the UK and 23 in Europe. Unfortunately, the

Crown Copyright © 2007

forecasts from the MOGREPS global ensemble are only available to a maximum lead time of T+36 due to a problem with the verification.

The forecast probability is calculated based on the number of ensemble members which forecast the event to occur. When an event threshold which coincides with an ensemble member forecast the probability of the event occuring is calculated as $p = k/(n+1)$ where n is the number of ensemble members and k is the number of members for which the event occurs. A continuous cumulative probability density function for event thresholds which fall between ensemble forecasts is calculated by linear interpolation. If an event threshold lies outside the spread of the ensemble then the probability distribution is smoothly reduced to zero, or increased to 1 as appropriate, over the 1 unit interval from the lowest or highest value ensemble forecast. The calculation of the forecast probability is illustrated in figure 10. For example, if an event threshold is defined at a value of 10 (knots, degrees or other unit), and the lowest forecast value from an ensemble with 9 members is 10.5, then the forecast probability for $P(x >= 10)$ is $0.95 = (n+0.5)/(n+1)$. The control forecast is included in the calculation of the forecast probability, which has been rounded to the nearest 10% to aid comparison between ensembles of different sizes. Nonetheless, some benefit is expected to remain for the ECMWF ensemble which has more than twice the number of ensemble members of MOGREPS.

Figure 11 shows that for forecasts of screen temperature greater than 10^oC the NAE ensemble is more skilful than the MOGREPS global ensemble, which in turn is more skilful than the ECMWF ensemble. The improved performance of the MOGREPS

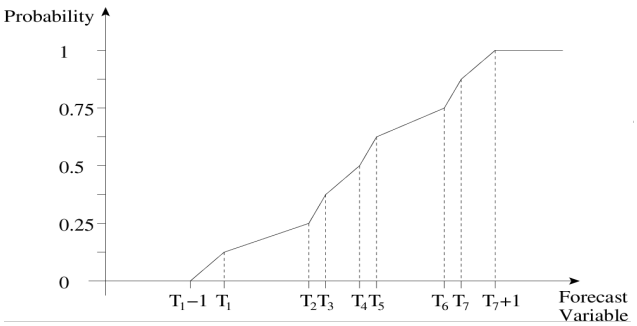

Figure 10. Illustration of how the forecast probability is calculated from an ensemble forecast with seven members.

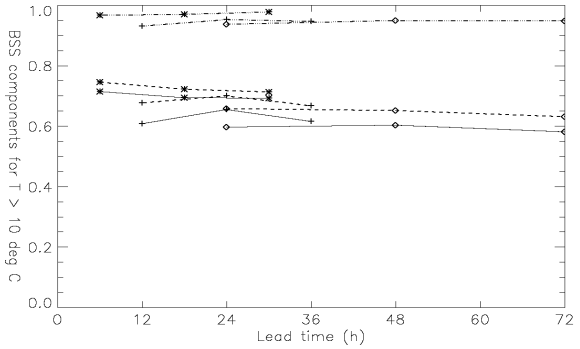

Figure 11. Brier skill score (solid), reliability (dot-dashed) and resolution (dashed) for T$> 10^o$C. Forecasts are shown the the MOGREPS NAE ensemble (stars), the MOGREPS global ensemble (pluses) and ECMWF ensemble (diamonds).

ensembles compared to the ECMWF ensemble is likely to be related to the use of the ETKF for specifying the initial condition perturbations. The ETKF is specifically designed for good performance in the short range, and these results appear to be consistent with that expectation. The better skill of the NAE ensemble, relative to the global ensemble, is likely to be due to its higher resolution. The ECMWF data are transferred to the Met Office on a 1.5^o grid, even though it is calculated at a much higher resolution than the MOGREPS global ensemble, and this may affect the scores seen here. A more detailed test would be required to fully assess the relative merits of the two systems.

Figure 12 shows the reliability diagram and sharpness for the same forecasts as Fig 11, for

Crown Copyright © 2007

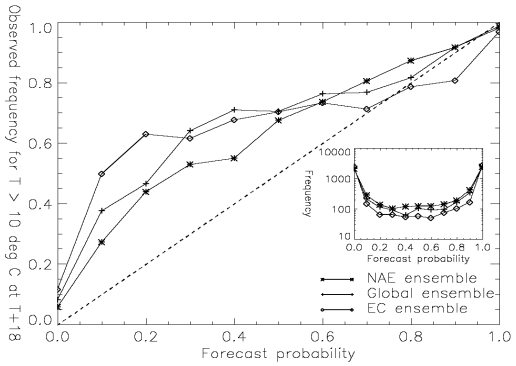

Figure 12. Reliability diagrams and sharpness for T$> 10^o$C. Forecasts are shown the the MOGREPS NAE ensemble (stars), the MOGREPS global ensemble (pluses) and ECMWF ensemble (diamonds).

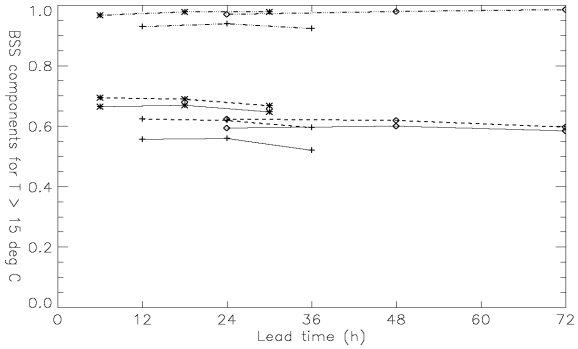

Figure 13. Brier skill score (solid), reliability (dot-dashed) and resolution (dashed) for T$> 15^o$C. Forecasts are shown the the MOGREPS NAE ensemble (stars), the MOGREPS global ensemble (pluses) and ECMWF ensemble (diamonds).

T+18 forecasts from the NAE ensemble and T+24 forecasts from the other two ensembles. Since the overall number of forecasts is the same, the sharpness of each forecast is easiest to see from the frequency with which each ensemble system forecasts the event to occur with an intermediate probability. One forecast is said to be more sharp than another if it has a greater propensity to forecast extreme probability (i.e. 0 or 1). Therefore, the ECMWF ensemble is the most sharp and the NAE the least, but it must be remembered that the ECMWF ensemble was the least skilful in this case. All the ensembles appear to be under-forecasting the temperature (all of the curves are mostly above the diagonal) but the NAE ensemble is the most reliable.

One known deficiency of the MOGREPS ensembles during the period verified is that they used values of the soil moisture which are derived from climatology. This means that for much of the period under consideration (March-May 2006) the soil moisture used by MOGREPS was too wet. This leads to a systematic under-estimation of the screen temperature and the bias seen in figure 12, particularly for the global ensemble. Since the deterministic models produce analyses of soil moisture the MOGREPS ensembles have been updated to use these values since the results reported here.

Verification for forecasts of screen temperature greater than T$> 15^o$C is shown in figures 13 and 14. From these we conclude that the NAE ensemble is the most skilful, but that the ECMWF ensemble is more skilful than the MOGREPS global ensemble. The most striking difference between the ECMWF and global ensembles is that the MOGREPS global ensemble is considerably less reliable than the other two ensembles. From figure 14 it is clear that the global ensemble has a considerable bias, which may be related to the soil moisture problem mentioned above.

The skill of forecasts for wind speeds of at least Beaufort force 5 (17 knots) are shown in figure 15. A major difference between this graph and the equivalent for temperature is that the resolution score of all three ensemble systems is very similar (and much lower than for forecasts of temperature). However, differences between the three systems still exist in

Crown Copyright © 2007

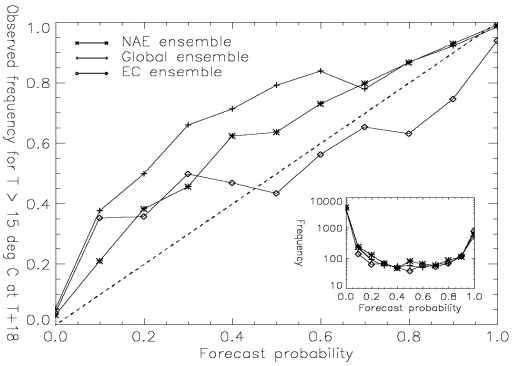

Figure 14. Reliability diagrams and sharpness for T$>15^oC$. Forecasts are shown for the MOGREPS NAE ensemble (stars), the MOGREPS global ensemble (pluses) and ECMWF ensemble (diamonds).

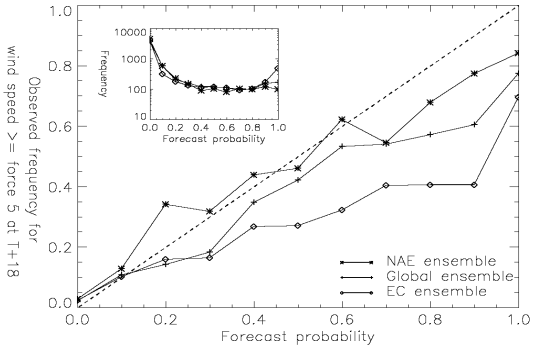

Figure 16. Reliability diagrams and sharpness for wind speed $>=$ force 5. Forecasts are shown the the MOGREPS NAE ensemble (stars), the MOGREPS global ensemble (pluses) and ECMWF ensemble (diamonds).

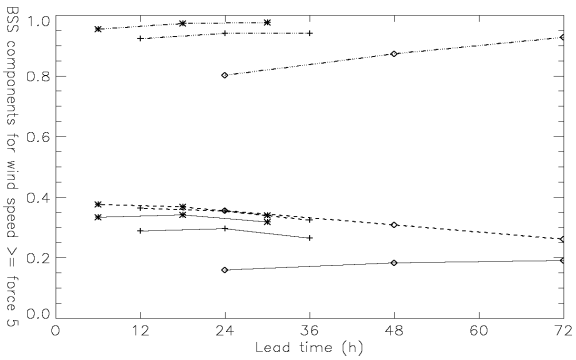

Figure 15. Brier skill score (solid), reliability (dot-dashed) and resolution (dashed) for wind speed $>=$ force 5. Forecasts are shown the the MOGREPS NAE ensemble (stars), the MOGREPS global ensemble (pluses) and ECMWF ensemble (diamonds).

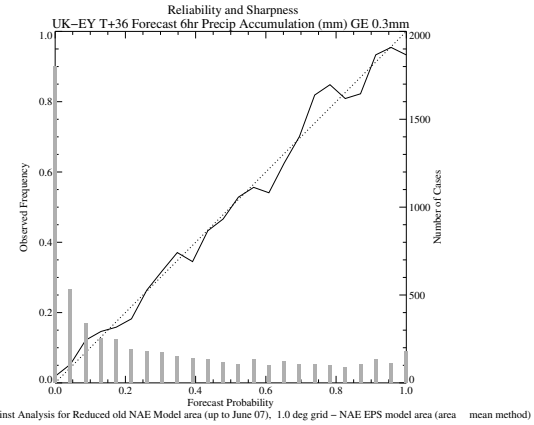

Figure 17. Reliability diagram and sharpness for NAE ensemble forecasts of the probability of 6hr rainfall greater than 0.3mm.

terms of the reliability of the forecasts. The NAE ensemble is the most reliable, followed by the global ensemble, with the ECMWF ensemble being the least reliable. From figure 16 it is seen that the reliability differences stem from an over-forecasting bias in the global and ECMWF ensembles.

The reliability of precipitation forecasts from the NAE ensemble is assessed in figures 17 and 18. The verification in these graphs is performed for forecasts verifying in the months of April to June 2006 against Nimrod radar-based rain analyses (Harrison *et al.*, 2000) aggregated to a 1^o resolution grid. The forecast probability for these data is calculated based on democratic counting (i.e. the probability is given by the fraction of members which forecast the event to occur). Figure 17 shows the verification of ensemble forecasts of the event '6-hour accumulated precipitation of at least 0.3mm'. The reliability curve is very close to the diagonal. This indicates that the combination of the initial condition perturbations with the stochastic physics scheme is very effective at creating spread in the precipitation forecasts. At the higher threshold of 5mm/6h (Fig 18) the reliability curve suffers from considerable noise due to sampling deficiency, but clearly lies below the diagonal. This suggests that the NAE ensemble is under-forecasting the frequency of moderate rain events. Nonetheless the ensemble clearly has some resolution.

Crown Copyright © 2007

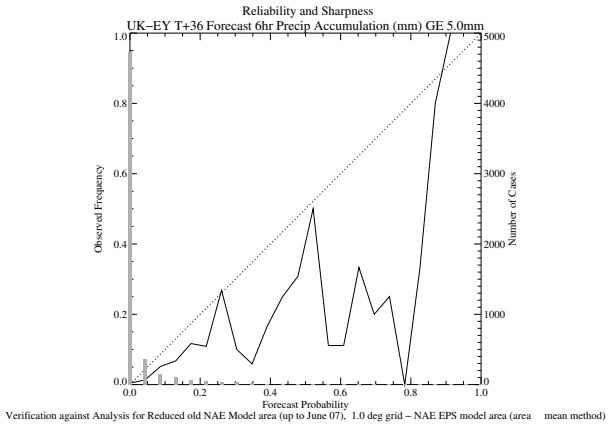

Figure 18. Reliability diagram and sharpness for NAE ensemble forecasts of the probability of 6hr rainfall greater than 5mm.

6 Conclusions

The Met Office Global and Regional Ensemble Prediction System (MOGREPS) has been running routinely since August 2005. Since that time it has been used by forecasters to provide guidance on the uncertainty in the weather for short-range forecasts. Subjective verification by Met Office forecasters has been very encouraging with numerous occasions of positive feedback about the usefulness of the forecast. This paper aims to support this assessment with an objective verification of the skill of the system.

A successful ensemble prediction system needs to simulate the effect of both initial and model-related uncertainties on forecast errors (Buizza *et al.*, 2005). In MOGREPS, the initial condition perturbations are generated using the Ensemble Transform Kalman Filter (ETKF). This is a computationally efficient version of the Ensemble Kalman Filter (EnKF) where the perturbations for each cycle are calculated as a linear combination of the forecast perturbations from the previous cycle. This is an appropriate technique for initial condition perturbations for a short-range ensemble, since the ETKF aims to quantify the uncertainties in the analysis. Two stochastic physics schemes, designed to represent the uncertainty due to model error (the stochastic convective vorticity for the global ensemble and the random parameters for both ensembles) are included. These schemes represent the uncertainties arising from the unresolved sub-grid-scales and deficiencies of the parameterised processes respectively. A complete description of the system can be found in part I.

The interaction between the initial condition perturbations and those coming from the stochastic physics schemes is an important aspect of the system. In order to analyse it, an experiment was set up in which perturbations were applied only to the initial condition (IC), only to the model errors (ME) or both (IC+ME). The spread in ME is approximately 13% of the spread due to the ICs and the spread for IC+ME experiment was around 5% larger than IC alone. After 72 hours all three ensembles (IC, ME and IC+ME) produce their largest spread over the same regions. This indicates that the growth and spatial distribution of spread is not exclusively related to the mechanism used to introduce the perturbations, but to the background flow.

Performance of MOGREPS has been assessed through routine verification over a number of months. The Brier skill score from the regional ensemble for forecasts of screen temperature greater than 10^{o}C and 15^{o}C, and wind speed greater than Beaufort force 5 is higher than that for the global or ECMWF

Crown Copyright © 2007

ensemble. This demonstrates that the ETKF perturbation method is at least competitive with singular vectors for short-range forecasts. It also shows the benefit from using high-resolution ensemble forecasts. Forecasts of 6-hour accumulated precipitation of more than 0.3mm from the NAE ensemble are reliable, but there is some indication of over-forecasting at higher rain accumulations.

The spread of the MOGREPS global ensemble broadly matches the error in the ensemble mean forecast. The spread of the ensemble in the tropics is smaller than the error of the ensemble mean, with the reverse being true close to the poles. Most of the energy of the initial condition perturbations is in the wind field, with less in the temperature perturbations. There is evidence that high spread in the ensemble is related to areas where largest errors are present in the forecast.

In future developments we hope to introduce localisation of the error covariances in the ETKF, as was suggested by Bowler (2006a). This should address the issues with the ensemble being over-spread near the poles. We also expect to introduce a model perturbation scheme based on kinetic energy backscatter (SKEB).

Acknowledgements

The authors would like to thank Rob Darvell and Anette van der Wal who provided the data which was used to generate some of the verification figures. The development of the MOGREPS system has involved the input of many people within the Met Office, and we would like to thank all those who have been involved.

References

Anderson, J.L., “A method for producing and evaluating probabilistic forecasts from ensemble model integrations”, *Journal of Climate*, **9**, 1518-1530 (1996)

Bowler, N.E., “Comparison of error breeding, singular vectors, random perturbations and ensemble Kalman filter perturbation strategies on a simple model”, *Tellus*, **58A**, 538-548 (2006a)

Bowler, N.E., “Explicitly accounting for observation error in categorical verification of forecasts”, *Monthly Weather Review*, **134**, 1600-1606 (2006b)

Bowler, N.E., Arribas, A., Mylne, K.R., and Robertson, K., “The MOGREPS short-range ensemble prediction system Part I: System description”, *Met Office technical report No. 497*, (2007)

Buizza, R., Houtekamer, P.L., Toth, Z., Pellerin, G., Wei, M., and Yuejian Z., “A comparison of the ECMWF, MSC and NCEP global ensemble prediction systems”, *Monthly Weather Review*, **133**, 1076-1097 (2005)

Hamill, T.M., and Colucci, S.J., “Verification of eta-RSM short-range ensemble forecasts” *Monthly Weather Review*, **125**, 1312-1327 (1997)

Hamill, T.M., and Juras J., “Measuring forecast skill: is it real skill or is it the varying climatology?”, *submitted to Quarterly Journal of the Royal Met. Soc*, (2007)

Harrison, D.L., Driscoll, S.J., and Kitchen, M., “Improving precipitation estimates from weather radar using quality control and correction techniques”, *Meteorol. Appl.*, **7**, 135-144 (2000)

Crown Copyright © 2007

Roulston, M.S., “A comparison of predictors of the error of weather forecasts” *Nonlinear Processes in Geophysics*, **12**, 1021-1032 (2005)

Tribbia, J.J., and Baumhefner, D.P., “Scale interactions and atmospheric predictability: An updated perspective”, *Monthly Weather Review*, **132**, 703-713 (2004)

Wilks, D.S., “Statistical methods in the atmospheric sciences (second edition)”, *Academic Press*, 627pp. (2006)

Crown Copyright © 2007